

the storage system as required. Such movement of data is often required during hardware failure and while performing preventive maintenance on the storage system (i.e.: swapping drives in RAID sets). The illustrative embodiments of the present invention provide duplicate copies of data which remain accessible through the VLUN even during the reconfiguration and repair of the RAID system and thus enable a user to access data at all times.

In one embodiment of the present invention, a host electronic device is interfaced with the network. A plurality of devices having storage mediums with device controllers, such as a RAID system, are also interfaced with the network. A software created VLUN is located on the host electronic device and acts as a virtual interface between a user of the host electronic device and the storage mediums. The users of the host electronic device, such as a file system, or database, send data write and read requests to the VLUN. The VLUN sends data write and read requests to a VLUN selected device controller for a device with a storage medium. The storage process beyond the VLUN is transparent to the user. The user has no direct interaction with the storage medium. In one implementation, the selected device is a RAID set, the device controller is a RAID volume controller, and the VLUN selects the RAID set based on a desired RAID level.

In another embodiment of the present invention, an apparatus interfaced with a network, includes a file system, processor and storage medium. The network is also interfaced with a plurality of storage devices which includes storage mediums. The apparatus further includes a software facility to create a virtual interface used to allocate data to available storage mediums. The storage mediums are transparent to the user of the apparatus and all read and write requests for storage are made through the virtual interface created by the software facility.

Description of the Drawings

Figure 1 is a block diagram of an environment suitable for practicing an illustrated embodiment of the present invention;

Figure 2 is a block diagram of the hierarchical location of the virtual interface (VLUN);

Figure 3 is a flow chart depicting the sequence of steps followed by an illustrative embodiment of the present invention to store data using a RAID volume controller;

Figure 4 is a flow chart of an alternate embodiment of the present invention indicating the steps followed upon the detection of an error in one side of a RAID set.

Figure 5 is a flow chart depicting the sequence of steps followed by an embodiment of the present invention to store data with the the VLUN directly allocating the data to storage mediums; and

Figure 6 is a flow chart of the sequence of steps followed by the embodiment of **Figure 5** in recovering from a detected error.

Detailed Description

The illustrative embodiments of the present invention provide a method for performing host based storage virtualization. A software created virtual interface, known a virtual logical unit (VLUN), is inserted between a user, such as a file system or database, on a host electronic device, and the end storage mediums located, in most cases, remotely over a network. The insertion of the VLUN abstracts the process of storing data such that the end storage mediums locations are opaque to the user. The user writes data to the VLUN and reads data from the VLUN. The VLUN is located on the host electronic device interfaced with the network. The location of the VLUN on the host electronic device enables the repair and reconfiguration of the storage mediums without requiring user knowledge of the details of the process. The local control of the VLUN allows a user to have uninterrupted access to the user's data even during reconfiguration and repair of the storage mediums.

Figure 1 depicts an environment suitable for practicing the illustrative embodiment of the present invention. A plurality of hosts, host A 2, host B 6, and host C 7 are interfaced with a network switch 8. Host A 2 includes a file system 1, memory 3, a VLUN 4 located in the memory, and a processor 5. Also interfaced with the network switch 8 are a plurality of storage devices 10, 14, 18, and 22. Each storage device 10, 14, 18, and 22 includes a volume controller 11, 15, 19, and 23, as well as a storage medium such as disk drives 12, 13, 16, 17, 20, 21, 24, and 25. The volume controller is used to write and read data from the storage mediums. Those storage mediums may be entire disk drives, or multiple volumes within a single disk drive. The

file system 1 located on host A 2 sends data to the VLUN 4. The VLUN 4 redirects the data to an available storage device 10, 14, 18, and 22 selected by the VLUN. The illustrated embodiment of the present invention may use a RAID system for backing up data. For instance, the VLUN 4 may send a collection of data for storage via the

5 network switch 8 to the RAID volume controller 15 located on storage device 14. The RAID volume controller 15 may then mirror the data on a mirror set made up of disk drives 16 and 17. "Mirroring" is a technique used to distribute multiple copies of the same data to two or more disks. Any time the data is updated, the data is changed for both copies. Similarly, the VLUN 4 may be used to read stored data. The VLUN 4 may

10 receive a read request from the file system 1 and send the request to the RAID volume controller 15. The RAID volume controller 15 retrieves the data from the disk drives 16 and 17 and sends the data to the VLUN 4 which sends it to the file system 1. The VLUN 4 provides an interface between a host user and a storage medium so that the process of storing or reading the data on the medium is hidden from the user. The

15 VLUN 4 may be moved between hosts 2, 6 and 7.

Figure 2 depicts a hierarchy utilized by the illustrative embodiment of the present invention. A file system 30 sends write and read requests to the VLUN 32. The file system 30 is a system that the operating system of the host 2 uses to organize files

20 and directories. The VLUN 32 is a software created virtual interface between the file system 30 and the rest of the network storage system. The specific locations of the storage devices (the logical units), which may be locally or remotely located, are opaque to the file system which reads and writes data through the VLUN 32. The VLUN 32 is capable of opaquely mapping data to many different types of storage

25 devices in response to one request from the file system 30. The VLUN 32 sends its own write and read requests to a RAID volume controller 34 selected by the VLUN. The RAID volume controller 34 writes and reads data on available disks 36 and 38 in response to the requests received from the VLUN. Each entity in the hierarchical diagram is able to give commands only to the entity directly underneath in the

30 hierarchical diagram. Thus the file system 30 may not write data directly to the RAID volume controller 34. Since the specific RAID volume controller location is immaterial to the file system 30 (which accesses the RAID volume controller 34 through the VLUN 32), the VLUN may substitute a different RAID volume controller (not shown) for the RAID volume controller 34 when storing data. Similarly, the disks 36 and 38

35 respond to a user's request for data by sending the data to the RAID volume controller 34 rather than directly to the VLUN 32 or file system 30. Those skilled in the art will